

CLAIM AMENDMENTS:

Listing of the Claims:

1. (Cancelled)
2. (Cancelled)
3. (Currently amended) A method according to Claim 2 40 wherein the positive anomaly examples are associated with fraud or software vulnerabilities.
4. (Currently amended) A method according to Claim 2 40 including developing the rule set using Higher-Order logic.
5. (Currently amended) A method according to Claim 4 including developing the rule set by:
 - a) forming an alphabet having selector functions allowing properties of the training data set to be extracted, together with ~~at least one of the following~~: additional concepts, background knowledge constant values and logical AND and OR functions,
 - b) forming current rules from combinations of items in the alphabet such that type consistency and variable consistency is preserved,
 - c) evaluating the current rules for adequacy of classification of the training data set in terms of indicating data anomalies in accordance with data flagging,
 - d) if no current rule adequately classifies the training data set, generating new rules by applying at least one genetic operator to the current rules, a genetic operator having one of the following functions: i) combining two rules to form a new rule, ii) modifying a single rule by deleting one of its conditions or adding a new condition to it, or iii) changing one of a rule's constant values for another of an appropriate type, and
 - e) designating the new rules as the current rules and iterating steps c) onwards until a current rule adequately classifies the training data set or a predetermined number of iterations is reached.
6. (Cancelled)
7. (Currently amended) A method according to Claim 6 40 for detecting

telecommunications or retail fraud from anomalous data.

8. (Previously presented) A method according to Claim 7 employing inductive logic programming to develop the rule set.
9. (Currently amended) A method according to Claim 8 wherein the at least one anomaly characterisation rule has a form that an anomaly is either detected or otherwise not detected by application of the rule according to whether or not a condition set of at least one condition associated with the rule is or is not fulfilled respectively.
10. (Cancelled)
11. (Currently amended) A method according to Claim 10 40 wherein a variable in the at least one anomaly characterisation rule which is defined as being in constant mode and is numerical is at least partly evaluated by providing a range of values for the variable, estimating an a respective accuracy for each value and selecting for the variable a at least one of the value values which has having optimum accuracy among the accuracies of the range of values.
12. (Currently amended) A method according to Claim 11 wherein the range of values is a first range with spacing between values which are relatively widely spaced, a single optimum accuracy value is obtained for the variable, and the method includes selecting a and relatively narrowly spaced second range of values in the optimum accuracy value's vicinity, the second range of values having narrow spacing relative to that of the first range, and the method also includes estimating an accuracy for each value in the second range and selecting a value in the second range having optimum accuracy.
13. (Currently amended) A method according to Claim 12 including filtering to remove rule duplicates and rule equivalents, i.e. any a rule equivalent being a rule having like but differently ordered conditions compared to another rule, and any or a rule which has conditions which are symmetric compared to those of another rule.
14. (Previously presented) A method according to Claim 13 including filtering to remove unnecessary 'less than or equal to' ('lteq') conditions.
15. (Previously presented) A method according to Claim 14 wherein the unnecessary lteq

conditions are associated with at least one of ends of intervals, multiple lteq predicates and equality condition and lteq duplication.

16. (Currently amended) A method according to Claim 8 including implementing an encoding length restriction to avoid overfitting noisy data by rejecting a rule refinement if the refinement has an encoding cost in number of bits which exceeds a cost of encoding positive examples covered by the refinement.
17. (Currently amended) A method according to Claim 8 including stopping construction of a rule in response to fulfilment of at least one of three stopping criteria, such criteria being:
 - a) the number of conditions in any rule in a beam of rules being processed is greater than or equal to a prearranged maximum rule length,
 - b) no negative examples are covered by a most significant rule, which is a rule that:
 - i) is present in a beam currently being or having been processed,
 - ii) is significant,
 - iii) has obtained a highest likelihood ratio statistic value found so far, and
 - iv) has obtained an accuracy value greater than a most general rule accuracy value, and
 - c) no refinements were produced which were eligible to enter the beam currently being processed in a most recent refinement processing step.
18. (Previously presented) A method according to Claim 17 including adding the most significant rule to a list of derived rules and removing positive examples covered by the most significant rule from the training data set.
19. (Previously presented) A method according to Claim 8 including:
 - a) selecting rules which have not met rule construction stopping criteria,
 - b) selecting a subset of refinements of the selected rules associated with accuracy estimate scores higher than those of other refinements of the selected rules, and
 - c) iterating a rule refinement, filtering and evaluation procedure to identify any refined rule usable to test data.
20. (Cancelled)

21. (Currently amended) Computer apparatus according to Claim 20 42 wherein the positive anomaly examples are associated with fraud or software vulnerabilities.
22. (Currently amended) Computer apparatus according to Claim 20 42 ~~programmed wherein the computer software programs the apparatus~~ to develop the rule set using Higher-Order logic.
23. (Currently amended) Computer apparatus according to Claim 22 ~~programmed wherein the computer software programs the apparatus~~ to develop the rule set by:
 - a) forming an alphabet having selector functions allowing properties of the training data set to be extracted, together with ~~at least one of the following:~~ additional concepts, background knowledge constant values and logical AND and OR functions,
 - b) forming current rules from combinations of items in the alphabet such that type consistency and variable consistency is preserved,
 - c) evaluating the current rules for adequacy of classification of the training data set in terms of indicating data anomalies in accordance with data flagging,
 - d) if no current rule adequately classifies the training data set, generating new rules by applying at least one genetic operator to the current rules, a genetic operator having one of the following functions: i) combining two rules to form a new rule, ii) modifying a single rule by deleting one of its conditions or adding a new condition to it, or iii) changing one of a rule's constant values for another of an appropriate type, and
 - e) designating the new rules as the current rules and iterating steps c) onwards until a current rule adequately classifies the training data set or a predetermined number of iterations is reached.
24. (Cancelled)
25. (Currently amended) Computer apparatus according to Claim 20 42 wherein the at least one anomaly characterisation rule has a form that an anomaly is either detected or otherwise not detected by application of such rule according to whether ~~or not~~ a condition set of at least one condition associated with that rule is or is not fulfilled respectively.

26. (Cancelled)

27. (Currently amended) Computer apparatus according to Claim 26 42 wherein a variable in the at least one anomaly characterisation rule is defined as being in constant mode and is numerical, and the computer apparatus is programmed to evaluate the at least one anomaly characterisation rule at least partly by providing a range of values for the variable, estimating an a respective accuracy for each value and selecting for the variable a at least one of the value values which has having optimum accuracy among the accuracies of the range of values.

28. (Currently amended) Computer apparatus according to Claim 25 programmed wherein the computer software programs the apparatus to filter out at least one of rule duplicates, rule equivalents and unnecessary 'less than or equal to' ('lteq') conditions.

29. (Currently amended) Computer apparatus according to Claim 25 programmed wherein the computer software programs the apparatus to stop construction of a rule in response to fulfilment of at least one of three stopping criteria, such criteria being:

- a) the number of conditions in any rule in a beam of rules being processed is greater than or equal to a prearranged maximum rule length,
- b) no negative examples are covered by a most significant rule, which is a rule that:
 - i) is present in a beam currently being or having been processed,
 - ii) is significant,
 - iii) has obtained a highest likelihood ratio statistic value found so far, and
 - iv) has obtained an accuracy value greater than a most general rule accuracy value, and
- c) no refinements were produced which were eligible to enter the beam currently being processed in a most recent refinement processing step.

30. (Cancelled)

31. (Currently amended) A computer software product according to Claim 30 43 wherein the positive anomaly examples are associated with fraud or software vulnerabilities.

32. (Currently amended) A computer software product according to Claim 30 43 wherein

the computer readable instructions provide for controlling computer apparatus to develop the rule set using Higher-Order logic.

33. (Currently amended) A computer software product according to Claim 32 wherein the computer readable instructions provide for controlling computer apparatus to develop the rule set by:

- a) forming an alphabet having selector functions allowing properties of the training data set to be extracted, together with ~~at least one of the following:~~ additional concepts, background knowledge constant values and logical AND and OR functions,
- b) forming current rules from combinations of items in the alphabet such that type consistency and variable consistency is preserved,
- c) evaluating the current rules for adequacy of classification of the training data set in terms of indicating data anomalies in accordance with data flagging,
- d) if no current rule adequately classifies the training data set, generating new rules by applying at least one genetic operator to the current rules, a genetic operator having one of the following functions: i) combining two rules to form a new rule, ii) modifying a single rule by deleting one of its conditions or adding a new condition to it, or iii) changing one of a rule's constant values for another of an appropriate type, and
- e) designating the new rules as the current rules and iterating steps c) onwards until a current rule adequately classifies the training data set or a predetermined number of iterations is reached.

34. (Cancelled)

35. (Currently amended) A computer software product according to Claim 30 43 wherein the at least one anomaly characterisation rule has a form that an anomaly is either detected or otherwise not detected by application of such rule according to whether or not a condition set of at least one condition associated with that rule is or is not fulfilled respectively.

36. (Cancelled)

37. (Currently amended) A computer software product according to Claim 36 43 wherein

the computer readable instructions provide for controlling computer apparatus to at least partly evaluate a variable in the at least one anomaly characterisation rule which is defined as being in constant mode and is numerical by providing a range of values for the variable, estimating ~~an a respective~~ accuracy for each value and selecting for the variable ~~a at least one of the value values which has having~~ optimum accuracy among the accuracies of the range of values.

38. (Currently amended) A computer software product according to Claim 35 wherein the computer readable instructions provide for controlling computer apparatus to filter out at least one of rule duplicates, rule equivalents and unnecessary 'less than or equal to' ('lteq') conditions.
39. (Previously presented) A computer software product according to Claim 35 wherein the computer readable instructions provide for controlling computer apparatus to stop construction of a rule in response to fulfilment of at least one of three stopping criteria, such criteria being:
 - a) the number of conditions in any rule in a beam of rules being processed is greater than or equal to a prearranged maximum rule length,
 - b) no negative examples are covered by a most significant rule, which is a rule that:
 - i) is present in a beam currently being or having been processed,
 - ii) is significant,
 - iii) has obtained a highest likelihood ratio statistic value found so far, and
 - iv) has obtained an accuracy value greater than a most general rule accuracy value, and
 - c) no refinements were produced which were eligible to enter the beam currently being processed in a most recent refinement processing step.
40. (New) An automated method of anomaly detection by applying a rule set to test for anomalies in data, the rule set comprising at least one anomaly characterisation rule, and the method incorporating the steps of:
 - a) providing a training data set incorporating positive and negative anomaly examples and expressed as digital data flagged to indicate presence and absence of data anomalies respectively,

- b) defining a rule generalisation based on logic of at least First-Order, and
- c) using computer apparatus to execute the steps of:
 - i) processing the rule generalisation to transform it into a more specific rule generalisation by adding at least one of a condition, a variable, a constant, a unification of variables and a function based on the training data set and relevant background knowledge consisting of at least one of concepts, facts of interest and functions for calculating values of interest,
 - ii) evaluating the more specific rule generalisation by applying it to the training data set to identify anomalies, and
 - iii) incorporating the more specific rule generalisation in the rule set if it classifies anomalies in the training data set adequately in terms of covering at least some of the positive anomaly examples, and
 - iv) applying the rule set to test data for anomaly detection therein, and
- d) providing an alert or a report to a user regarding anomaly detection in the test data resulting from operation of the method.

41. (New) A method according to Claim 40 wherein after the step of evaluating the more specific rule generalisation, the computer apparatus is used to iterate the processing and evaluating steps for one or more rules in the more specific rule generalisation which do not classify anomalies in the training data set adequately.

42. (New) Computer apparatus for anomaly detection by applying a rule set to test for anomalies in data, the rule set comprising at least one anomaly characterisation rule, the computer apparatus being associated with a training data set incorporating positive and negative anomaly examples and expressed as digital data flagged to indicate presence and absence of data anomalies respectively, a rule generalisation being defined based on logic of at least First-Order, and the computer apparatus incorporating computer software which programs it to execute the steps of:

- a) processing the rule generalisation to refine it and render it more specific by adding at least one of a condition, a variable, a constant, a unification of variables and a function based on the training data set and relevant background knowledge consisting of at least one of concepts, facts of interest and functions

for calculating values of interest,

- b) evaluating the more specific rule generalisation by applying it to the training data set to identify anomalies, and
- c) incorporating the more specific rule generalisation in the rule set if it classifies anomalies in the training data set adequately in terms of covering at least some of the positive anomaly examples,
- d) applying the rule set to test data for anomaly detection therein, and
- e) providing an alert or a report to a user regarding anomaly detection in the test data resulting from operation of the computer apparatus.

43. (New) A computer software product comprising a computer readable hardware medium containing computer readable instructions for controlling operation of computer apparatus to implement anomaly detection by applying a rule set to test for anomalies in data, the rule set comprising at least one anomaly characterisation rule, the computer readable instructions being associated with a training data set incorporating positive and negative anomaly examples and expressed as digital data flagged to indicate presence and absence of data anomalies respectively, a rule generalisation being defined based on logic of at least First-Order, and the computer readable instructions providing a means for controlling the computer apparatus to execute the steps of:

- a) processing the rule generalisation to refine it and render it more specific by adding at least one of a condition, a variable, a constant, a unification of variables and a function based on the training data set and relevant background knowledge consisting of at least one of concepts, facts of interest and functions for calculating values of interest,
- b) evaluating the more specific rule generalisation by applying it to the training data set to identify anomalies, and
- c) incorporating the more specific rule generalisation in the rule set if it classifies anomalies in the training data set adequately in terms of covering at least some of the positive anomaly examples,
- d) applying the rule set to test data for anomaly detection therein, and
- e) providing an alert or a report to a user regarding anomaly detection in the test data resulting from operation of the computer apparatus.